

SUPERMICRO STORAGE SOLUTIONS FOR HIGH-PERFORMANCE COMPUTING

No-Compromise Storage Solutions to Speed Up HPC Workloads



TABLE OF CONTENTS

Introduction
HPC Storage Requirements2
$Software-Defined-Storage \ Architectural \ Choices \ \dots \ 6$
Supermicro HPC Storage Reference Architecture $\ldots \ldots 6$
All Flash Tier Storage Systems8
All Flash Data Platform Software10
Capacity Tier Storage Systems
Conclusion 14
For More Information15

Introduction

HPC storage systems must handle massive volumes of data and deliver extreme performance to support demanding workloads, such as large-scale simulations, data-intensive scientific computations, and machine- learning applications. The data generated by these applications can be unstructured, including images, videos, and text, or structured, such as numerical data and databases.

HPC storage systems must be highly scalable to accommodate growing data sets and highly available, incorporating robust fault tolerance and disaster recovery mechanisms. These systems must handle parallel access from multiple nodes with low latency and high throughput.



HPC Storage Requirements

The specific requirements for HPC storage vary depending on the application. For example, scientific simulations may require high-performance storage for large datasets and rapid data transfer. Machine learning applications may require fast access to training data and efficient storage for model checkpoints. In all cases, HPC storage systems must deliver the performance and capacity needed.

High-performance computing workloads need multiple storage tiers to balance performance needs and costs.

- An All-Flash storage tier to cache input for data-hungry application servers and deliver high bandwidth continuously to meet demand.
- A Storage Capacity tier containing all current and historical data, with ample storage capacity.

As the value of transforming data into actionable insights becomes increasingly clear, the drive to collect and process data intensifies. The more detailed and comprehensive the data, the more trends can be identified and leveraged for a competitive advantage. Companies are turning to Supermicro to harness the transformative potential of various HPC and data-intensive applications and workloads.

HPC customer and workload storage challenges: HPC customers face unique storage challenges. The sheer volume of data generated and processed by HPC applications demands robust, scalable, and high-performance storage solutions.

Key Challenges:

- Data Volume and Velocity:
 - Massive Data Growth: HPC applications often generate enormous datasets, requiring significant storage capacity.
 - Real-time Data Processing: Data's rapid generation and consumption necessitate low-latency storage systems.
- Performance and Scalability:
 - High-Throughput and Low-Latency: HPC workloads demand high-performance storage to keep up with intensive computations.
 - Scalability: As HPC workloads grow, storage systems must scale seamlessly to accommodate increasing data volumes and performance needs.
- Data Integrity and Security:
 - Data Protection: data must be protected from loss, corruption, and unauthorized access.
 - Data Recovery: Robust data protection and recovery mechanisms are essential to minimize downtime
- Cost-Effectiveness:
 - Balancing Performance and Cost: HPC storage solutions must be optimized to deliver high performance while minimizing costs.
- Integration and Management:



- Heterogeneous Environments: HPC systems often involve diverse hardware and software components, requiring seamless integration.
- Management: Managing large-scale, high-performance storage systems can be challenging.

Addressing the Challenges: To address these challenges effectively, a Parallel file system with a Tiered storage architecture is recommended:

Parallel File Systems



Parallel file systems are designed to meet the demands of high-performance computing (HPC) environments. They provide efficient access to large datasets and enable parallel I/O operations across multiple nodes. These file systems typically utilize a distributed architecture, storing data on multiple servers or storage devices to achieve parallel I/O capabilities.

A key characteristic is the ability to provide high-performance access to large datasets through parallel I/O operations, where multiple nodes can simultaneously read or write data from the file system. This significantly improves performance, especially for large-scale data transfers and intensive computational workloads. Key characteristics of a Parallel File System are

- Metadata Server: The server keeps track of file information like location, size, and ownership.
- Large Storage capacity: Data is stored on multiple Petabyte-sized Storage Servers.
- Parallel I/O: Multiple clients can simultaneously read or write data from different parts of a file, significantly increasing throughput.
- Global Namespace: Users can access data through a single, unified namespace, regardless of where the data is physically stored.
- High-Speed Networking: Parallel file systems often utilize high-bandwidth network connections, such as InfiniBand, to optimize data transfer.



Advantages of Parallel file systems for HPC environments.

- Scalability, allowing the file system to grow as the number of nodes and data volumes increases.
- High availability, with redundant data placement and fault tolerance mechanisms to ensure data integrity and minimize downtime.

Organizations can choose between Open-Source or Commercial Parallel file systems depending on their specific needs:

Open-Source Parallel File Systems

- Lustre: One of the most widely used open-source parallel file systems, known for its scalability, high performance, and flexibility.
- BeeGFS: A high-performance, easy-to-use parallel file system with excellent scalability and performance.

Commercial Parallel File Systems

- WEKA: A modern parallel file system that leverages NVMe storage to deliver high performance and low latency.
- VAST Data Platform: accelerates HPC by delivering the performance and scale of parallel file systems with the simplicity of NAS, all at archive economics.
- DDN Exascaler: high-performance parallel file system designed to accelerate AI, analytics, and HPC applications.
- Hammerspace Data Platform: a parallel global file system with automated data placement that supports files and objects.

These are just a few of the many parallel file systems available today. The best choice for a particular application will depend on various factors, including performance requirements, scalability needs, and budget constraints.

Challenges with Parallel File Systems:

- Complexity: Managing and optimizing parallel file systems can be challenging due to their distributed nature.
- Data consistency: Ensuring data integrity when multiple clients access the same file concurrently.
- Cost: Deploying a large-scale parallel file system can require significant hardware investment.
- Support: specialized expertise is needed, only available with selected providers.

Tiered Storage Architecture

Tiered storage architecture is a strategy that involves organizing storage into different tiers based on performance, capacity, and cost. This approach optimizes storage utilization and performance by matching data to the appropriate storage tier.

Characteristics and Features

- Tiered Storage Levels: Tiered storage systems typically consist of multiple tiers, such as high-performance SSDs, high-capacity HDDs, and cloud storage.
- Data Migration or Hierarchical Storage Management: Data is automatically migrated between tiers based on access patterns and defined policies.
- Caching: High-performance tiers can cache frequently accessed data to improve performance.
- Compression and Deduplication: Data compression and deduplication techniques can reduce storage requirements and improve performance.



Benefits of Tiered Storage

- Improved Performance: Placing frequently accessed data on high-performance tiers can significantly improve overall system performance.
- Reduced Storage Costs: Less frequently accessed data can be stored on lower-cost, higher-capacity tiers, reducing overall storage costs.
- Enhanced Data Protection: Data can be replicated or backed up across multiple tiers to ensure data durability and disaster recovery.
- Scalability: Tiered storage systems can be easily scaled to meet growing data storage needs.

Challenges of Tiered Storage

- Complexity: Managing multiple tiers and implementing data migration policies can be complex.
- Performance Overhead: Data migration and caching can introduce performance overhead, especially during peak usage periods.
- Cost: Implementing a tiered storage system can involve significant upfront costs, including hardware and software.
- Integration with Existing Systems: Integrating tiered storage into existing IT infrastructure can be challenging, especially for legacy systems.

Organizations can optimize storage performance, reduce costs, and improve management efficiency by carefully designing and managing a tiered storage architecture.

- High-Performance Tier:
 - Purpose: Stores active data and provides low-latency, high-throughput access for critical workloads.
 - Technologies: High-speed SSDs, NVMe storage, and high-performance file systems.
- High-Capacity Tier:
 - Purpose: Stores historical data and less frequently accessed files.
 - \circ ~ Technologies: Capacity storage, cloud storage, and tape storage.

Key Considerations:

- Data Lifecycle Management: Implement policies to automatically migrate data between tiers based on usage patterns and retention requirements.
- Data Protection and Recovery: Employ robust data protection strategies, including backups, snapshots, and replication.
- Performance Optimization: Optimize storage systems for specific workloads, such as data analytics, machine learning, and simulation.
- Security: Enforce strong security measures to protect sensitive data from unauthorized access and cyber threats.
- Monitoring and Management: Utilize advanced monitoring and management tools to track system performance, identify bottlenecks, and proactively address issues.

By carefully designing and managing tiered storage architecture, HPC organizations can effectively address their storage challenges and unlock the full potential of their high-performance computing capabilities.







Tiered storage architectures typically use cost effective storage for the High-capacity Tier, using a storage system that stores large amounts of unstructured data like images, videos, and documents each with unique identifiers and metadata, within a flat data structure, allowing for easy access and retrieval without needing a hierarchical file system like traditional storage methods; essentially, it's a way to store data without needing to organize it into folders or directories, making it ideal for large volumes of diverse data.

Software-Defined Storage (SDS) Architecture Choices

Software-defined storage offers flexibility by separating storage software from hardware. Two key architecture choices are Share Nothing and Shared NSPOF (No Single Point of Failure) Hardware.

- Share Nothing features independent nodes with dedicated resources. SDS distributes data (sharding) across nodes, which are managed by a distributed software layer that communicates over a network. It uses standard hardware with local storage. Advantages include high horizontal scalability, node-level fault tolerance, good local performance, cost-effectiveness, and no single hardware failure points. Disadvantages are management complexity, network latency for remote data, data redistribution overhead, and challenges with consistency. Examples include Ceph and Cassandra.
- Shared NSPOF Hardware utilizes shared, redundant storage enclosures managed by SDS. The software pools storage, provisions volumes, and manages data placement. Compute nodes connect via a storage fabric. Advantages include simpler management, data mobility, potentially lower local network overhead, efficient resource utilization, and hardware-level high availability. Disadvantages include limited scalability compared to shared nothing, potential bottlenecks in shared components, higher initial cost, and potential single points of failure at the SDS control or fabric level. Examples include virtualized SANs.

Choose the architecture that best meets your needs for scalability, performance needs, budget, and management preferences. Share nothing suits massive scale-out, while shared NSPOF offers simpler management and hardware-level resilience. Both provide viable SDS solutions.

Supermicro HPC Storage Reference Architecture

Supermicro offers a broad server and storage product line, including high-performance, all-flash storage ideal for active data and high-capacity systems with spinning disks for high-capacity, low-cost capacity storage.

Supermicro's engineering has been focused on addressing data-intensive computing challenges for many years. It has developed a general scale-out storage architecture that serves the needs of HPC workloads such as manufacturing, life sciences, energy exploration, research, and financial services. Each deployment will have specific needs, and Supermicro's solutions team will work with you and use its expertise to customize and design the best solution for your needs.



Lustre PFS Architecture



Many customers use the Supermicro storage systems for high-performance applications, including autonomous driving, largescale modeling and analytics, biomedical research, and data mining. Nowhere are the demands more intensive than in HPC. Efficiently managing massive data set transfers on tiered storage systems running parallel file systems requires an uninterrupted stream of bi-directional data movement for modeling, simulation, and analytics workloads on powerful datahungry rack-scale computing systems.

Supermicro offers a scale-out storage reference architecture to address the needs described above. The architecture is flexible and scalable, able to adapt to a wide range of storage-intensive application needs and grow as data sets expand. It consists of two storage tiers connected to the HPC application servers via a high-speed network.

All-Flash Performance Tier

The core of the solution, the all-flash tier, holds the data actively used by applications. Typically, this tier accounts for 10 to 20 percent of overall data and utilizes all-flash storage servers to store data on a distributed, scale-out file system for various clients. It can also coordinate data tiering to the capacity-storage tier and the cloud. Using servers optimized for storage instead of more costly purpose-built appliances enables linear scaling, increasing bandwidth and the capacity to store more data as additional servers are added to the cluster. Storage clusters in this tier must be interconnected with the highest bandwidth available networking, either 400 Gbps Ethernet or 400 Gb/s InfiniBand, to maintain the lowest possible latency. Supermicro offers several options for All-Flash performance Tier, including a 1U Petascale server and a 2U High Availability Storage Server.

Capacity Storage Tier

While the all-flash tier acts as a cache for the organization's most actively used data, it can't just be deleted when it is no longer frequently used. It still has value and needs to be retained in storage optimized for capacity and durability instead of



performance, which reduces the overhead for 80 to 90 percent of data in this category. The servers to support this mission have a high density of disk drives and comparatively low bandwidth, as well as IOPS requirements for handling data transfers in and out of the storage cluster. Similarly, the requirements for networking are relaxed, and 100 Gbps Ethernet is an economical choice that balances cost and performance for capacity storage needs. Supermicro offers several options for Capacity Tier Storage, including High Density 60/90 Hard Disk drive-based Top Loading Storage, a 24 HDD 2U SimplyDouble storage server, and, where needed, a U.2 SSD-based Hyper Rack Mount Server.

Ideally, software to manage this tier needs to provide a searchable, protected content repository for unstructured data where it participates in the global namespace provided by the all-flash tier. It needs to focus on protecting data from loss due to hardware failures and threats, including ransomware or even accidental deletion. It needs to manage storage tiering to migrate data to cloud-based storage when directed by data lifecycle policies.

The software must provide many of the same capabilities as the all-flash tier, but focus on storage lifecycle management. It needs to scale non-disruptively and continues to operate without data loss due to the failure of any of the tiers' components. However, it must also protect and preserve data through erasure coding, data protection, versioning, remote replication, object locking, threat mitigation, and implementation of lifecycle data retention and migration policies.



Supermicro Reference Architecture for HPC Storage

All-Flash Tier Storage Systems



All-Flash Tier based on Supermicro Petascale 1U and 2U Storage Systems:

Key Features

- Support high core count latest dual Intel socket and single AMD single socket
- Support 1U E1.S and E3.S and 2U E3.3 TLC, QLC and CXL device and up to 1PB all flash in 2U
- Optimized thermal design with EDSFF Spec
 Balanced PCIe lanes for front SSD and rear IO
- design
- Up to 30M 4KB Random Read IOPS and 230+ GB/s 128KB Sequential Read BW.

1U Petascale All-Flash Storage Specifications

	1U E1.S 24-Bay	1U E3.S 16-Bay	1U E3.S 16-Bay	1U E3.S 16-Bay
Model Number	SSG-121E-NES24R	SSG-121E-NE316R/SSG-121E- NEX316R	ASG-1115S-NE316R/ASG-1115S- NEX316R	SSG-122B-NE316R
CPU	Dual 4 th Gen Intel [®] Xeon [®] Scalable Processors up to 270W TDP	Dual 4 th Gen Intel [®] Xeon [®] Scalable Processors up to 270W TDP	Single AMD EPYC TM 9004, 9005 Series Processor up to 300W TDP	Intel® Xeon® 6700 / 6500 series processors up to 270W
Memory	32 DIMM slots; Up to 8TB DDR5- 4800	32 DIMM slots; Up to 8TB DDR5- 4800	24 DIMM slots; Up to 6TB DDR5-4800	32 DIMM slots; Up to 8TB DDR5-5200
Expansion	2 PCIe 5.0 x16 FHHL slots 2 PCIe 5.0 x16 AIOM slots	2 PCIe 5.0 x16 FHHL slots 2 PCIe 5.0 x16 AIOM slots	2 PCIe 5.0 x16 FHHL slots 2 PCIe 5.0 x16 AIOM slots	2 PCIe 5.0 x16 FHHL slots 2 PCIe 5.0 x16 AIOM slots
Storage	24 hot-swap 9.5mm/15mm E1.S	16 hot-swap 7.5mm E3.S Up to 4 E3.S CXL x8 slots (optional)	16 hot-swap 7.5mm E3.S Up to 4 E3.S CXL x8 slots (optional)	16 hot-swap 7.5mm E3.S drives
Networking	via AIOM	via AIOM	via AIOM	via AIOM
Power Supplies	Redundant 2000W Titanium level	Redundant 1600W Titanium level	Redundant 1600W Titanium level	Redundant 2000W Titanium level
Workloads	Software-defined Storage, Data Intensive HPC, NVMe Over Fabrics Solution, In-Memory Computing, Private & Hybrid Cloud			



All-Flash Tier based on Supermicro HA 2U all-flash array

Key Features

- Transactional Scale-up Storage
- Failover clustering (Active-Active)
- High Performance NVMe
- Tiered scale-up by adding JBOD



	2U U.2 24-Bay
Model Number	SSG-221E-DN2R24R
CPU (per node)	Single S^{th} Gen $Intel^*$ Xeon * Scalable Processors up to 350W TDP
Memory (per node)	8 DIMM slots; Up to 2TB DDR5-5600
Expansion (per node)	2 PCIe 5.0 x8 FHHL slots 2 PCIe 5.0 x16 FHHL slots
Storage	24 hot-swap 15mm dual port U.2
Networking (per node)	Dual 10GbE BaseT with Intel [®] X710
Power Supplies	Redundant 2000W Titanium level
Workloads	Software-defined Storage, Mission Critical Database, Parallel File System Metadata node

High-Speed Networking

The capacity and performance of the servers' NVMe storage must be matched with appropriate network connectivity. Each server supports multiple NVIDIA ConnectX-7 SmartNICs, each with a single NVIDIA InfiniBand NDR 400 Gb/s port. These occupy each server's PCIe 5.0 x16 expansion slots. A pair of NDR InfiniBand switches (available as part of SSE-MQM9700-NS2F), each with 64 nonblocking ports, supports two independent data fabrics.

All Flash Data Platform Software

Supermicro partners with multiple software-defined storage vendors, enabling customers to choose based on their specific requirements and needs.

- WEKA Data Platform: WEKA is architecting a new approach to the enterprise data stack built for the AI era. The WEKA® Data Platform sets the standard for AI infrastructure with a cloud and AI-native software solution that provides seamless data portability across on-premises, cloud, and edge environments. It transforms legacy data silos into data pipelines that make GPUs, AI, and other performance-intensive workloads run faster and more efficiently, delivering microsecond latency performance at scale. WEKA helps the world's most innovative enterprises and research organizations accelerate time to market, discovery, and insights with AI, including 12 of the Fortune 50.
- VAST Data Platform: offers a defined storage solution specifically designed for high-performance computing (HPC) workloads, with a unique architecture that combines high-capacity flash storage with advanced data reduction techniques, enabling users to store and access massive datasets with exceptional performance. This simplifies data management and reduces operational overhead. VAST's solution is known for its scalability, high performance, and



ease of use, making it an attractive option for organizations with demanding HPC needs. Supports a wide range of data types and protocols, eliminating the need for separate storage silos; seamlessly scales to meet the ever-growing demands of modern data centers, performance for both reads and writes, accelerating data-intensive applications. It also has an intuitive interface and automation features that simplify data management and reduce operational overhead.

- **DDN Exascaler:** Provides high-performance storage solutions designed for demanding HPC workloads. Key features include: Massively Scalable Architecture, which enables seamless scaling to accommodate massive datasets and thousands of concurrent users; delivers exceptional throughput and low latency; and maximizes the utilization of expensive computing resources. In addition, it incorporates advanced data reduction techniques to optimize storage utilization and reduce costs, and offers comprehensive data protection features, including data replication and disaster recovery. Finally, it provides intuitive management tools and automation capabilities to simplify operations.
- **Hammerspace:** Hammerspace offers a unique approach to HPC storage, providing a global data platform that unifies unstructured data across various locations, including on-premises, cloud, and edge environments. Key features include a single global namespace that simplifies data management and access, high-performance access to data regardless of location, and data orchestration capabilities that automate data placement and movement. This eliminates data silos and enables seamless data sharing and collaboration across distributed teams and applications. Hammerspace's solution is particularly well-suited for organizations with complex data workflows and the need for high performance and scalability.
- **Qumulo**: Qumulo provides a modern, scale-out file storage solution that seamlessly integrates storage capabilities, offering a unified platform for diverse unstructured data. This approach allows organizations to manage data within a single, scalable namespace, simplifying data management and reducing complexity. Qumulo Cloud Data Fabric unifies edge, data center, and cloud into a single cohesive data storage platform, enabling organizations to make data available to compute and GPU at any location, at any time, instantly. It provides high-performance file access through intelligent caching and a high-performance distributed file system, delivering the benefits of both file and object storage, whether in the data center or in the cloud. It supports NFS, SMB, and S3, efficiently handling various workloads, from high-performance computing to cost-effective archiving, with consistent data services. Qumulo on Supermicro expands customer choice from all-NVMe for the most demanding use cases to hybrid-NVMe for enterprise workloads that need high capacity and performance. Qumulo's software-only approach ensures full functionality and optimal performance on any supported Supermicro appliance.



Capacity Storage Systems

4U Cloud Scale Top Loading Storage System



Key Features

- Single, Twin and JBOD configurations
- Up to 2.7 PB capacity in 90 bay or 1.8 PB capacity in 60 bay system. (w/ 30TB HDD)
- Enterprise-grade serviceability
- Best-in-Class performance

X14 Single-Node Top Loading Storage Specifications

1 000000000000000000000000000000000000	DADADAD P
	9

4U Single-Node 60-Bay



4U Single-Node 90-Bay

Model Number	SSG-542B-E1CR60	SSG-542B-E1CR90/
CPU (per node)	Intel® Xeon® 6700 / 6500 series processors up to 350W	Intel® Xeon® 6700 / 6500 series processors up to 350W
Memory (per node)	16 DIMM slots; Up to 2TB DDR5-6400 / 5200	16 DIMM slots; Up to 2TB DDR5-6400 / 5200
Expansion (per node)	1 PCIe 5.0 x8 (in x8) HHHL slot 3 PCIe 5.0 x16 (in x16) HHHL slots	1 PCle 5.0 x8 (in x8) HHHL slot 3 PCle 5.0 x16 (in x16) HHHL slots
Storage (per node)	60 hot-swap 3.5" SAS/SATA 4 rear hot-swap 2.5" NVMe	90 hot-swap 3.5" SAS/SATA 4 rear hot-swap 2.5" NVMe
Storage Controller	Broadcom 3916 or 3816	Broadcom 3916 or 3816
Power Supplies	Redundant 2000W Titanium level	Redundant 2700W Titanium level
Workloads	Big Data & Analytics, Data Lake, Telco & Cloud Service Providers, Contents Repositories, HPC and AI/ML Workloads, Financial Service & Healthcare Image Archives, Government Data Protection	



2U Hyper Servers for Capacity Storage

Key Features

- Best-in-class computing performance for the most demanding workloads
- Higher memory bandwidth and capacity with DDR5 RDIMM and MRDIMM
- Direct connect NVMe on PCIe Gen 5 : higher throughput and lower latency
- Flexible expansion slot and networking configuration



CPU	Dual and Single Intel® Xeon® 6900/6700/6500 Series Processors and AMD EPYC [™] 9005/9004 Series Processors
Memory	32 DIMM (Dual Xeon [®] 6700/6500), 24 DIMM (Dual Xeon [®] 6900, Dual/Single AMD EPYC [™] 9005/9004), 16 DIMM (Single Xeon [®] 6700/6500), 12 DIMM (Single Xeon [®] 6900, Single AMD EPYC [™] 9005/9004)
Expansion	8 PCIe 5.0 x8 (or 4 PCIe 5.0 x16) slots in 2U Up to 2 AIOM* slots for flexible networking options
Storage	NVMe SSD support with up to 24 U.2 drives in 2U (Optional SAS and SATA configurations)
Power Supply	
Workloads	SW defined Storage, Parallel File System

Not all data needs to be in the fastest, all-flash tier. Indeed, for a large amount of big data to be processed and analyzed, it can be retained on a cost-effective storage capacity tier until it is needed again.

The Data Platform Software must support migrating data to a second storage tier. It interoperates with several scale-out distributed storage platforms that can be used to support this tier. The large data sets used in HPC workloads will be stored on the capacity tier and migrated to the performance tier as needed.

Capacity Storage Tier Software

There are many choices for capacity tier software that interoperates with the Supermicro architecture, including commercial and open-source options. This object store software can be deployed on All-Flash storage servers for higher performance.

Quantum ActiveScale: Quantum ActiveScale object storage supports the Capacity Storage tier. This durable storage repository connects to the all-flash tier and supports active archiving and long-term data retention. The software provides an unlimited scale with support for exabytes of capacity. Newly added nodes will join the existing cluster in servicing requests, resulting in unlimited performance as well. ActiveScale achieves high performance through highly parallelized software that load balances data and operations across cluster nodes. Data access is always available with rolling system upgrades and the capability to tolerate component and site failures. The software is also optimized for low management overhead. For example, the software will automatically detect failed disks and repair any affected data. The replacement of failed disks can be postponed and bundled together, resulting in less frequent data center visits and involvement from IT staff.



OS Nexus QuantaStor: provides a software-defined storage solution that supports object storage alongside block and file storage. Its Storage Grid Technology dramatically reduces the complexity of cluster management by combining storage appliances across public and private clouds to manage all components as one and enabling seamless integration with a wide range of applications and tools. 1 It leverages Ceph technology for high performance and scalability, offering robust data protection and management capabilities such as encryption, replication, and snapshots. It can be deployed on various hardware platforms, providing flexibility and cost-effectiveness for organizations of all sizes. Storage Grid Technology: All major storage protocols, including NFS/SMB, iSCSI/FC, and S3, are supported along with end-to-end security.

Cloudian: HyperStore is a software-defined, scale-out storage platform known for its strong S3 API compatibility, enabling seamless integration with cloud-native applications. It allows users to build on-premises or hybrid cloud storage solutions, managing vast amounts of unstructured data with exabyte scalability. Key features include robust data protection through replication and erasure coding, advanced security with encryption and WORM capabilities for ransomware protection, and multi-tenancy support. HyperStore offers a cost-effective alternative to traditional storage, running on standard hardware and providing unified file and Object management for diverse workloads.

Minio: a high-performance, software-defined storage server built for cloud-native and edge deployments. Architected with a focus on simplicity, scalability, and Amazon S3 compatibility, Minio allows users to build their private cloud storage infrastructure on commodity hardware. Its lightweight and container-friendly design makes it ideal for modern application development, data analytics, and AI/ML workloads requiring fast and reliable access to unstructured data. Minio's distributed architecture enables seamless horizontal scaling to handle massive datasets, while its strong consistency model ensures data integrity. Furthermore, its open-source nature and vibrant community contribute to its rapid development and broad adoption as a versatile storage software solution.

Conclusion

Supermicro delivers complete end-to-end Rackscale Solutions for Business Advantage.

Choosing the proper storage solution for your HPC Workloads is critical and significantly impacts business outcomes for all workloads, especially High Performance Computing requirements. Selecting storage architecture by working with experts from Supermicro is a strategic business decision; it takes advantage of the complete end-to-end Rack Scale solution capabilities, including server, storage, networking, and management tools. Supermicro can customize the solution for sufficient bandwidth and low latency required to meet the application and business requirements, and provide deployment and maintenance services.

With Supermicro's unique building block approach, capacity, performance, and cost can be balanced to deliver a solution that meets application needs within budget constraints. Supermicro's state-of-the-art manufacturing facilities in the USA ensure a Secure Supply Chain, short fulfillment cycles, and the organization's deployment and time-to-value milestones.





For More Information:

Supermicro Storage Systems: <u>www.supermicro.com/storage</u>

Supermicro Storage Solutions: <u>https://www.supermicro.com/en/solutions/software-defined-storage</u>

Supermicro HPC Solutions: <u>www.supermicro.com/hpc</u>

Contact us at <u>www.supermicro.com/en/contact</u>

